

The Day the Record Clears: Statutory Expiry of Criminal-Record Disclosure and Ex-Offenders' Return to Work

Bolero

Abstract

Pre-data manuscript scaffold. Registered methods and planned displays are prepared; all empirical results and evidence-dependent conclusions remain unfilled.

Keywords: Registered study, results pending

PRE-DATA DRAFT: empirical placeholders are not findings.

Title

The Day the Record Clears: Statutory Expiry of Criminal-Record Disclosure and Ex-Offenders' Return to Work

Abstract

A conviction may impede employment through information shown to employers, persistent characteristics of the person, or both. This results-unfilled manuscript carries forward an existing preregistration that uses scheduled expiry of Denmark's private criminal-record certificate as an individual event-time design. The source proposes single-observed-conviction decision cohorts from 2009-2017, employee job starts from BFL, and sanction-specific clocks of 24 and 36 months from the final decision and 60 months from release, conditional on a historical legal encoding that remains unfinished. H1 alone is confirmatory. A never-disclosed traffic-fine battery is estimated before treated-arm results; its comparison with the fine arm and a within-cohort placebo constrain causal interpretation after the required estimates exist. The main estimate is `[[RESULT:abstract_main]]` and its registered interpretation is `[[RESULT:abstract_claim]]`. These placeholders also admit nonactivation, pre-data unresolved requirements and operational failure. Schedule, industry, employed-worker mobility, job-quality and benefit-state analyses remain exploratory. Their outcomes

cannot raise the confirmatory claim ceiling. This is source-PROTOCOL conversion, not replication of an already published empirical paper; no empirical findings exist in this manuscript.

Keywords

Registered study; results pending.

Introduction

The employment consequences of a conviction can outlast the sentence. The source protocol asks whether an employer-facing record helps sustain those consequences. A record can contain information that affects hiring while also marking accumulated disadvantage or persistent differences. Those accounts imply different effects of a change in disclosure status, even when they produce similar cross-sectional employment gaps. The question concerns the scheduled information event, not the total effect of conviction or incarceration.

The original preregistration positions its contribution relative to petition-based expungement, automatic record remediation and employer audit studies. Petitioning can select people whose circumstances are already changing; automatic remediation removes that particular choice, but the salience and institutions of disclosure differ across settings. Prescott and Starr’s Michigan expungement study and Agan and colleagues’ automatic-remediation work are therefore distinct reference points for the source question. The source also draws on callback experiments, research on hiding record questions, Nordic incarceration studies and legal scholarship on certificates. These literatures motivate the question without supplying a new Danish treatment effect, a numerical prior or a basis for changing the registered analyses. The inherited literature-search record, including failed queries and bibliographic version differences, is retained in docs/sources/literature.md. This conversion makes no exhaustive novelty claim.

The source design centers each person on a mechanically scheduled private-certificate disclosure expiry. It proposes that different sanction horizons and a legally never-disclosed comparison class make timing patterns informative about the disclosure interpretation. That argument is conditional: the historical legal table, the exact offence list and the traffic-sanction screen must be frozen before data access. Present-day police guidance is not enough to validate those historical definitions. Employer certificate requests are unobserved, and neither schema documentation nor the statutory schedule shows that a person or prospective employer actually obtained a new certificate at that date.

The confirmatory question is correspondingly narrow. Among people non-employed at the start of a month, does the probability of starting an employee job change across scheduled clearing? The source permits a causal ceiling only for that scheduled-date quantity under maintained assumptions and successful claim-gate adjudication. Attribution specifically to certificate clearing over residual disclosed-arm timing institutions is framed as strongly consistent with the evidence. Employer-information attribution is at most consistent-with because workers may change search

timing around the same known date. No prediction for other countries, certificate regimes or self-employment is licensed.

Source registration and hypotheses

The source is prereg.yaml version 1.0, created 21 August 2026, frozen at commit aa65f6764ddf before source-reported microdata access. This conversion is dated 5 September 2026, with the precise provenance timestamp in docs/replication/source-provenance.yaml. The original protocol and this manuscript are not an already-published empirical study. No independent reviewer has accepted this conversion, and historical self-audit findings do not constitute acceptance. The format-only Transcriber will preserve authored scientific content; the host owns LaTeX/PDF compilation.

H1 predicts a positive change in monthly employee-job-finding hazard after scheduled clearing, without a comparable within-cohort placebo-anniversary jump. H0 is the source substantive-null interpretation: a smooth hazard through the date would be compatible with persistent characteristics rather than exclusion caused at this margin. H0 is not an additional test; a nonsignificant H1 is bounded by its confidence interval, never accepted as a true zero or proof of the alternative mechanism. The exact frozen statements and source forecasts are preserved in the owning Atom protocols.

H2 predicts that proportional hiring changes are greater in the frozen certificate-intensive industry group. H3 predicts migration of the contrast across the sanction-specific schedule: F versus T at 24 months, and S at 36 versus 24 months, with the custodial release/decision-anchor comparison as exploratory evidence. H2, H3 and H3c are exploratory in the final source registration. Additional exploratory questions concern hours, wages, annual wage income, match durability, the three-way public/private certificate comparison, benefit-state composition and employed-worker mobility.

Four explicit source amendments must remain visible. A1 replaced the earlier mandated project stop after low simulated H3a power with a demote-and-narrow family iteration, leaving H1 alone. A2 changed the placebo significance screen from alpha 0.05 to 0.01. This reduces false blocks under a true-null placebo but allows some real anomalies to pass that would have failed the old screen; both the power motive and that cost were disclosed. The original-screen verdict is a mandatory sensitivity. A3 added the traffic disqualification/custodial-sanction exclusion and A4 added case/court fields for appeal resolution and linkage. All were made before registration/data access but were not preregistered in the earlier binding design. A1 remains subject to the source editorial-adjudication condition; no acceptance is inferred here.

Data and population

The present boundary permits planning with all curated FSV registers but reports no validated tables. The source's historical NOTE discusses a smaller upload list, which does not constrain this run and does not establish current data access. Required current candidates are KRAF for decisions, KRIN for releases, BFL for employee jobs, BEF and source-compatible population/spine information, an unresolved emigration mapping involving VNDS_UD/CPST, IND for annual wage income

and OF for benefit states. Exact fields, Danish documentation locators, reference semantics, join cardinalities and gaps appear in docs/protocol/data-contract.md. FSV's PNR metadata does not prove an available crosswalk to the source PERSON_ID/AJO_PERSON_ID pseudonyms.

The proposed population comprises decisions of guilt with sanction during 2009-2017 with no earlier source-qualifying conviction observed since 1980. This is not necessarily a lifetime single-conviction cohort. The source requires age 18-55 at the clock origin, Danish residence at index and population-spine presence at event-window start. The fine arm F uses source sanction codes 211-225 for the Criminal Code plus an exact drugs list that remains to be supplied. S uses fully suspended 121-124; C uses unconditional 111, 114, 117 with a linkable release; T uses Road Traffic Act fine codes with A3 exclusions for 219, disqualification or a custodial element. Other source-excluded sanctions, unclassifiable offence codes and unresolved appeals enter the sample-flow exclusions. No omitted code list is inferred from a present-day website.

Appealed decisions require the same-JOURNR final-instance decision date after ordering by the historically valid AFG_AFGRETKO instance map and date. That rank map is unfinished. C uses the earliest release strictly after index and within eight years, first matching person plus case number and then the source person-level fall-back. Under the source contingency, a release-linkage rate below 90% removes C from pooled confirmatory H1, with F+S documented before treated contrasts; C results remain exploratory. Expected sample sizes in the source are planning assumptions. The realized sample and exclusion flow are [[RESULT:sample_flow]]; cohort composition is [[RESULT:cohort_description]]; pooled scope is [[RESULT:analysis_scope]].

Time zero is the month of scheduled clearing under the verified historical table: source F/T pseudo+24 months, S+36 months from the final decision, C+60 months from actual release. The source asks for event months-24 through+18, at least 12 observable pre months and six post months. Long custodial releases may move the window beyond available follow-up; documentation does not guarantee observability. Widening to 2018-2019 is only the source contingency, not an automatic conversion choice.

A person-month enters the risk set if no BFL compressed job covers its first day, using the source inclusive start<=first day<=end rule. The outcome is any recorded job start in that month. Multiple nonemployment spells are permitted and person is the clustering unit. When several starts occur, attributes are taken from the start with most paid hours. Censoring occurs in the month of reconviction, emigration, death, age 60 or coverage end. Source-silent ties, open job ends, left-censored spell duration, missing birth dates and exact spine/death/emigration mappings remain pre-data requirements. They are neither imputed nor silently solved by changing the cohort.

Identification and analysis methods

The source estimator is a stacked discrete-time hazard LPM on the non-employed risk set. Monthly event-bin indicators-24 through+18 use-1 as reference, with calendar-year-by-month and nonemployment-spell-duration fixed effects (months 1-36 separately, 37+pooled) and arm effects only in the pooled specification. The

implementation is R/fixest with person-cluster CR 1 covariance; a complete identical Stata fallback remains a source-required deposit. There are no primary person effects or other controls. Identification rests on locally smooth counterfactual job-finding hazards and valid source measurement, not on covariate adjustment alone.

The main contrast is the equal-bin mean of estimated coefficients in $[0, +5]$ minus $[-6, -1]$; the within-cohort placebo is $[-12, -7]$ minus $[-18, -13]$. These are model-coefficient contrasts with their covariance, not substitutions of raw pooled-window means. The sole confirmatory test is two-sided Wald alpha 0.05, with a positive prediction. A significant negative estimate is disconfirming. Raw hazard levels remain descriptive companions, distinct from coefficient estimates.

Stage B first produces the seven source checks, sample flow and composition displays without hypothesis-relevant event models. Coverage must align with published benchmark definitions within 5%; CPR linkage must reach 95%; the source baseline employment/hazard ranges are 30%-80% and 1%-15%; key variables must reach 99% nonmissing, industry 95%, and unresolved appeals no more than 5%. Decision-date heaping, job-start date quality, boundary composition and the mechanical calendar-month density check have the exact source thresholds and dispositions. Reconciliation, scoped warnings and unresolved stops are `[[RESULT:quality_checks]]`. The checks do not prove legal correctness, local smoothness or identifying variation. All quantities remain unobserved in this draft.

The T battery is estimated and archived first, including its full profile and one source-specified non-gating 90% hazard-ratio band. The exact band construction is unfinished and remains a pre-data gate; no new estimator or fixed equivalence margin fills it. H1 and F are then estimated. Only then can the T claim kill compare a positive significant T jump (two-sided 0.05) with 50% of F's jump. The within-cohort kill requires a positive significant placebo (two-sided 0.01) with magnitude at least half the main contrast. Both kills withdraw causal interpretation while leaving all results reportable. The support rule separately requires $\text{abs}(z_placebo) < 2.576$ and $\text{abs}(\text{delta_placebo}) < 0.5 \text{ delta_main}$. Failure of a support clause without the explicit kill pattern means non-support, not automatically design failure.

Source planning calibration: placebo-kill detection is 0.472 at a true placebo jump of half the floor, while T-kill detection is 0.372 and 0.950 at true T jumps of 50% and 100% of F. Source true-null false-fire rates are 0.0017 and 0.0242. These are historical simulation values, not achieved empirical power; gate survival does not exclude near-threshold contamination. Any evidence-dependent gate claim below must repeat these detection limits. The mandatory old-screen sensitivity uses 1.96 and is shown beside the A2 screen without calling A2 a tightening.

The source H1 module also retains its conditional donut, person-FE, matched-T, constant-composition, heaped-date exclusion, composition-reweighted and reconviction-reset estimates, plus exploratory KOEN/age heterogeneity. The primary estimate stays unchanged when a variant is triggered. Anticipation requires a significantly lower-3:-1 hazard than-6:-4 and adds the+1:+6 versus-6:-3 donut; a reset estimate differing by more than one primary SE requires both estimates in main text and reconviction-free scope. A date-heaping sensitivity that changes sign or loses significance is foregrounded alongside the original. Exact source-silent weighting, trigger

measurement and related implementation choices remain unfilled, with all variants still inventoried.

Exploratory analyses

H3a is F24 minus T24, with independent-arm variance; H3b is S36 minus S24, with within-person covariance. H3c is the C release+60 contrast; E7 preserves the paired release+60 and decision+60 estimates. H3a/H3b/H3c/E7 stay in the source schedule Holm family. The exact E7 test-unit counting and duplicate-release handling are unfinished; neither dropping a member nor inventing a joint test is permitted. A release-only jump is the source schedule-consistent reading and jumps at both anchors are anniversary-suspect. Nonsignificance at an anchor is not demonstrated absence. The source's low planning power means migration nonsignificance is uninformative rather than evidence against the schedule.

H2 compares proportional cause-specific hazard changes: `delta_cert` divided by its preperiod hazard minus the analogous other-industry ratio. The frozen DB07 list is 64-66, 80, 84, 85, 88.91 and 49.32; the source's transport label refers operationally to taxi 49.32. Bootstrap inference uses 999 person resamples and percentile 95% intervals. Fewer than 1000 certificate-intensive hires or an insufficient required disclosure cell prevents theta estimation and permits only the source-compliant descriptive composition report. E4 refines the grouping to private-certificate private screening sectors, public-certificate sectors including 84, and low-screening sectors, predicting `private>low>=public`. It retains division 84 exclusion and the source-triggered missing-industry sensitivity. Exact historical authority-sector mapping, adjusted-interval decisions and the source F4 model-to-hazard-level display mapping remain unfinished; F4 cannot substitute raw means for missing model output; the audit's proposed additional child-sector exclusion is not added.

E6 concerns people employed in month-1 and estimates workplace mobility plus its certificate-intensive destination interaction. It is independently reportable from H2 because the population and behavior differ. Source AJO_ARBNR_SENR's documented coverage ends in 2017, so a full-window identity mapping remains unresolved. The source mechanism Holm family spans H2, E4, E6. Multiple component tests and bootstrap p/adjusted-CI definitions are not silently collapsed into a new family.

E1 compares first-full-month hours and log narrow wage among pre/post hires. E2 compares within-person annual wage income in clearing-year+1, +2 versus -2, -1. E3 compares the share of pre/post matches lasting at least 180 days. Wage/income winsorization uses the source within-year 1 st/99 thpercentiles; hours>400 are missing, with the additional source phrase winsorized hours explicitly unresolved. These outcomes remain distinct within the quality family {E1, E2, E3}; missing annual income or censored job ends do not invent zeros or erase independent components.

E5 classifies month-start nonemployment as registered unemployment/activation versus other using OF intervention spells and estimates both source H1-style contrasts and their difference. The exact historical class/overlap table is unfinished. E5 is the final YAML's singleton uncorrected audit test; E7 appears alongside it only in presentation. This audit may describe a benefit-clock pattern but cannot independently establish or exclude confounding or rescue a failed H1 gate.

Results

Sample, measurement and descriptive profile

Table 1 reports `[[RESULT:sample_flow]]`, including source appeal/sanction exclusions and C linkage/observability losses. Table 2 reports `[[RESULT:cohort_description]]`. Table 3 reports `[[RESULT:quality_checks]]`, with every threshold, written reconciliation, scope decision and unresolved requirement. Figure 5 shows the scheduled-date calendar-month distribution, `[[RESULT:clock_density]]`, and Figure 6 the boundary composition audit, `[[RESULT:boundary_composition]]`. If shared prerequisites remain pending or fail, this section instead states the exact pre-data or measured boundary and lists downstream nonactivation; no empirical conclusion is substituted.

Figure 3 shows descriptive monthly employee-job-finding hazards by arm: `[[RESULT:hazard_profiles]]`. These are at-risk means, not causal effects or adjusted event coefficients. A missing arm or bin is labeled with its registered reason. Source wording anticipating a flat T panel is not a finding.

Confirmatory estimate and falsification evidence

Figure 1 presents the pooled coefficients and 95% intervals, `[[RESULT:pooled_event_profile]]`, and Figure 2 shows arm-specific coefficients, `[[RESULT:arm_event_profiles]]`. Table 4 reports `[[RESULT:h1_main]]`, the within-cohort placebo `[[RESULT:h1_placebo]]`, and the selected F+S+C/F+S scope. Table 5 reports the prior archived T battery `[[RESULT:traffic_battery]]` and its single non-gating interval `[[RESULT:traffic_band]]`. The post-estimation claim gates and mandatory A2/original-screen comparison are `[[RESULT:claim_gates]]`. Source calibration is placebo-kill detection 0.472 at half-floor and T-kill detection 0.372/0.950 at 50%/100% of F, with false-fire 0.0017/0.0242; these planning rates accompany any interpretation invoking gate survival.

Summary (H1). `[[RESULT:h1_summary]]` If the registered state is `supported_stable` or `supported_sensitive`, the text begins “Consistent with our hypothesis” and reports the positive scheduled-date contrast within the maintained assumptions and narrow ceiling. `supported_sensitive` additionally states that the original-screen verdict differs. If negative, it begins “Contrary to our hypothesis” and reports disconfirmation. If either kill fires, it states that causal interpretation is withdrawn and retains the main sign and interval. If `not_supported`, it identifies the failed support clause without calling it a kill. If bounded, it reports no evidence of an increase and the 95% upper limit `[[RESULT:h1_bound]]`, without accepting the null or proving persistent traits. If operationally failed, unactivated or source-unresolved, it states precisely that no registered main conclusion exists.

Table 10 reports `[[RESULT:robustness_results]]`. Triggered but unavailable variants remain listed with reasons; no primary estimate is hidden. The anticipation, heaping, composition, comparator and reconviction-reset restrictions that apply to the prose are `[[RESULT:robustness_scope]]`. The exploratory KOEN/age results or unresolved specification status are `[[RESULT:heterogeneity]]`.

Exploratory analyses

The schedule contrasts in Table 6 are [[RESULT:schedule_results]]. The release/decision-anchor reading is [[RESULT:custodial_anchors]]. When H3a/H3b are positive and source-Holm significant, this section calls them exploratory corroboration. When either is nonsignificant, it states that migration is uninformative at the source calibration and gives the confidence bounds. Opposite directions, mixed anchors and unavailable C components are reported individually. These results never upgrade the H1 verdict.

Table 7 and Figure 4 report the industry and public-certificate patterns, [[RESULT:industry_results]] and [[RESULT:industry_profiles]]. If H2's reporting gate prevents theta estimation, they show the permitted descriptive composition and the exact threshold reason; E4 receives its own disposition. A significant negative theta is disconfirming of the predicted industry signature; nondetection weakens localization to statutory screening sectors without accepting equality. The already-employed workplace-mobility and destination-interaction results in Table 7 are [[RESULT:mobility_results]]. They are interpreted separately from the non-employed hire population, with unresolved workplace coverage disclosed.

Table 8 reports hours/log wages, annual wages and match survival separately, [[RESULT:quality_results]]. Discordant quality measures are not averaged into a composite. Table 9 reports the benefit-state contrasts and difference, [[RESULT:benefit_results]]. Neither the quality module nor the benefit audit licenses a new confirmatory or causal mechanism claim. All exploratory family decisions and any unfinished correction are [[RESULT:exploratory_inference]].

Discussion

The central source question concerns an effect of reaching a scheduled disclosure date. The interpretation justified by the complete registered evidence is [[RESULT:discussion_primary]]. A supported main result under maintained assumptions can carry the source causal ceiling for this date-specific quantity, with strong consistency rather than definitive exclusion of disclosed-arm timing artifacts. Gate-calibration limitations remain explicit: placebo-kill detection 0.472 at half-floor and T-kill detection 0.372/0.950 at 50%/100% of F, source true-null false-fire 0.0017/0.0242. A gate failure, negative result, nonsignificant result or unresolved execution receives its prespecified wording instead of a favorable narrative.

The industry and employed-mobility evidence bears on a narrower mechanism claim, [[RESULT:discussion_mechanism]]. Even aligned signatures cannot isolate an employer-belief effect from worker search timing. If the industry signature is absent or opposite, employer localization is weakened and worker-side retiming remains undischarged. If the custodial anchors are both associated with changes, the source anniversary concern is foregrounded. Mixed and incomplete evidence limits scope component by component; nothing exploratory changes the sole confirmatory family.

A null result has a useful but limited meaning only through its interval and assumptions. The source bound describes the largest compatible increase, not an equivalence test and not evidence that records never affect hiring. Small effects, stale certificates,

monthly date imprecision, selection and incorrect historical coding can remain compatible with imprecision or attenuation. Conversely, the source’s assertion that misdating only attenuates is not a general guarantee established by these checks. The source’s planning floor is not an observed smallest meaningful effect or a test of absence.

The protocol itself has unfinished requirements: historical legal vintages and offence lists, instance ranking, key/spine/censoring maps, left-censored spell coding, the exact T band, selected exploratory model details and multiplicity treatment. The original audit is evidence for those concerns, not permission for this conversion to repair them. The current all-FSV boundary does not eliminate these gaps. Restrictions arising from source gaps, operational failures, missing components and data coverage are [[RESULT:discussion_limits]]. No conditional conclusion should be read as current evidence while its result slot remains unfilled.

The scope remains Danish employee-job finding around the specified disclosure schedule in the registered population. Self-employment transitions, actual certificate requests and an isolated employer-versus-worker causal decomposition are not observed outcomes here. Implications for policy must remain at this margin and obey the main verdict and all active restrictions. The conditional conclusion is [[RESULT:conclusion]].

Ethics, data access and reproducibility

The source proposes secondary analysis of pseudonymized registers within DST Forskningsservice, contingent on the relevant approval and source-required pre-IPA deposits. This conversion makes no new approval claim and contacts no individuals or institutions. No rows or computed outcome summaries were accessed. No private protocol was uploaded for public lookup. Only schema/documentation and the authorized local source text were inspected.

Future output release follows the source floors of at least five persons, with at least three workplaces and no workplace above 75% for workplace/industry cells. Exact current clearance and collapse mappings remain to be established; no potentially disclosive cutpoints or individual listings enter the manuscript. The six implemented figure templates consume only cleared plot-ready aggregate tables and show explicit unavailable-input reasons. Synthetic interface checks occur outside the World and never supply Results placeholders. Native protocols, the complete membership roster and source-conversion audit remain authoritative; this paper, its figures and eventual PDFs are derived presentations.

References and source note

Prescott, J.J. and Starr, S.B.(2020). Expungement of Criminal Convictions: An Empirical Study. *Harvard Law Review* 133(8), 2460-2555. [Publisher record](#). The source YAML also cites a 2019 version; this manuscript uses the verified publication year.

Agan, A.Y., Garin, A., Koustas, D.K., Mas, A. and Yang, C. S.(2024). Can You Erase the Mark of a Criminal Record? Labor Market Impacts of Criminal Record Remediation. NBER Working Paper 32394, DOI 10.3386/w32394. [Metadata record](#).

Search metadata verified; direct page retrieval was blocked, so no new full-text findings are claimed.

Danish Police. Afgørelser på din straffeattest. [Current guidance](#), consulted 5 September 2026. Present-day information is not validation of the study’s historical period.

Additional source-cited positioning, retained without a new full-text verification claim: Agan and Starr (2018), Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment, QJE 133(1), 191-235; Ahmed and Lång (2017), The Employability of Ex-Offenders, IZA Journal of Labor Policy 6(6); Bhuller, Dahl, Løken and Mogstad (2020), Incarceration, Recidivism, and Employment, JPE 128(4), 1269-1324; Blumstein and Nakamura (2009), Redemption in the Presence of Widespread Criminal Background Checks, Criminology 47(2), 327-359; Doleac and Hansen (2020), The Unintended Consequences of Ban the Box, Journal of Labor Economics 38(2), 321-374; Selbin, McCrary and Epstein (2018), Unmarked?, Journal of Criminal Law and Criminology 108(1). Full inherited search/nearest-work inventory and access limits are in docs/sources/literature.md; it includes Onal 2022 and the source’s Danish/Nordic legal and institutional references without extending the analysis.

Original scientific registration: The Day the Record Clears, prereg.yaml v1.0, 21 August 2026, commit aa65f6764ddf; private local source identified by hashes in docs/replication/source-provenance.yaml. This manuscript was prepared by the source-protocol graph author on 5 September 2026. No empirical outcome or independent acceptance is asserted.




Fig. 1 Source F1: pooled H1 event-bin coefficients and registered 95% intervals in percentage points relative to month -1; scheduled month 0 and within-cohort placebo windows shown. Actual profile: `[[RESULT:pooled_event_profile]]`. Use the registered F+S+C or contingent F+S scope; missing input is disclosed, never filled.

Arm-specific event studies
Registered display; reviewed results not yet available.

Fig. 2 Source F2: F/S/C/T source coefficients and 95% intervals on a common scale. F and T pseudo dates use 24 months, S36, C release60 under source historical mapping. T is the archived battery, not an assumed flat series. `[[RESULT:arm_event_profiles]]`.

Descriptive job-finding hazards
Registered display; reviewed results not yet available.

Fig. 3 Source F3: raw monthly employee-job-finding hazards among those non-employed at month start, separately by arm. Levels are descriptive and no confidence interval is invented. `[[RESULT:hazard_profiles]]`.

Industry cause-specific hiring profiles
Registered display; reviewed results not yet available.

Fig. 4 Source F4: certificate-intensive versus other and E4 three-way source model hazard profiles, with fractions displayed as percentages. The exact source model-to-level mapping remains a pre-data requirement; missing model output cannot be replaced by raw means or reconstructed from normalized coefficients. `[[RESULT:industry_profiles]]`. Coarse class/cell/reporting gates and unavailability reasons remain visible.

Scheduled-clearing month density
Registered display; reviewed results not yet available.

Fig. 5 Source F5: calendar-month share of scheduled dates by arm against the source 2xuniform line (16.7% rounded display). The plot is descriptive; exact source threshold-rounding convention must be frozen before gate execution. `[[RESULT:clock_density]]`.

Boundary risk-set composition audit
Registered display; reviewed results not yet available.

Fig. 6 Source F6: magnitude of signed standardized differences for the source KOEN and five-year-age categories in the risk set before versus after clearing, with 0.10 reference. `[[RESULT:boundary_composition]]`. Larger differences trigger the source reported reweighting sensitivity, not deletion of the main result.